WHITE PAPER

11 MYTHS LEAVING YOU VULNERABLE TO MODERN ATTACKS

Distributed denial-of-service (DDoS) attacks have risen dramatically in their size, scale, distribution, and sophistication in recent years, which is highlighted by some record-breaking attacks. Unfortunately, many organizations still hang on to some outdated thinking about how to defend themselves — assuming that their defenses are sufficient, or worse, that they're unlikely to be a target. The truth is: Victims of these attacks span all key industries, from financial services to ecommerce to gaming. In fact, attacks on critical public infrastructure, including healthcare, energy and utilities, education, and transportation have been of particular concern. In 2023, Akamai protected a customer in the Asia-Pacific region from a massive 900 gigabits per second (Gbps) attack. Later in the same year, Akamai prevented a 634 Gbps, 55 million packets per second (Mpps) attack that featured a complex mix of attack vectors — one of the largest attacks ever against a U.S. financial services customer. This is on top of the largest DDoS attack Akamai has mitigated to date: a 1.44 Tbps, 385 Mpps globally distributed attack lasting nearly two hours. These events make it clear that cybercriminals continue to target critical pillars of the economy.

Although the scale of these attacks may lead some smaller organizations to believe their risk of becoming a target of a DDoS attack is low, the reality is that business-critical services and applications in every industry are easy targets. The rise of politically and ideologically motivated hacktivists and the relatively low cost of DDoS as a service offered by cybercriminal groups, such as Killnet and Anonymous Sudan, have made nearly everyone a possible target. It's not just the initial attack that organizations need to worry about, either. DDoS attacks are increasingly being used as a smokescreen to distract network and security resources while attackers attempt simultaneous ransomware DDoS attacks (RDDoS) or other nefarious exploits like triple extortion campaigns. Finally, the increasing and alarming adoption of artificial intelligence tools to orchestrate highly sophisticated and distributed DDoS attacks creates a significant defensive challenge for businesses and public institutions that need to ensure consistent availability and performance.

As threats become more complex and evolve almost by the day, many myths about DDoS protection unfortunately still exist — some of them even encouraged by security vendors. DDoS protection must be a key tenet of any security strategy, so understanding the danger that these myths pose are critical to your DDoS defense.

MYTH 1

Total capacity indicates the full extent of mitigation resources available

Although total capacity is important, a simple network capacity number can be misleading by leaving out important details. Organizations that are evaluating DDoS protection technology solutions need to ask:

- How much network capacity is dedicated to consuming attack traffic?
- How many of the mitigation system's resources are **explicitly dedicated** to stopping attacks?
- How many of the network and system resources are available to deliver clean traffic to all customer origins on that platform and to each unique tenant?

These questions are critical because if the total network capacity includes other requirements, such as content delivery, the actual DDoS defense capacity might be only a fraction of what the provider is claiming.

DDoS defense capacity isn't just limited to technology, either. At some point, if the technology stops working effectively, will there be dedicated human resources for escalations, incident response, and fine-tuning mitigation? The most robust mitigation combines automation and machine intelligence with human expertise to offer in-depth protection.

Tip

Look deeper into the differences between a provider's total network capacity and its platform stability, as well as how much capacity it has for attack mitigation and clean traffic delivery. They should be considered unique segments. For example, capacity should be dedicated by purpose, such as network routing of attack traffic, stopping or mitigating attack traffic, and delivering clean traffic back to the data center.

MYTH 2

DDoS protection from internet service providers and/or cloud service providers are sufficient

Unfortunately, many organizations still think that the protection offered by their internet service provider (ISP) is all they need. The truth is: ISPs typically only provide retooled, commercial, off-the-shelf DDoS protection with limited bandwidth. Their hardware is shared between their own infrastructure and yours, which means constrained capacity and CPU cycles. DDoS attacks now are so massive that they can overwhelm both infrastructures, and the ISPs will null route (or blackhole) your traffic to prevent collateral damage to other production resources. By blackholing all traffic, businesses lose legitimate traffic and services from end users, thereby making the attack successful by taking business offline for all practical purposes.

Additionally, while cloud service providers (CSPs) often allow customers to set their own controls and maintain sovereignty over their security posture within the CSP's cloud environment, most of the CSPs themselves typically reject any accountability and end up charging customers for the illegitimate DDoS traffic. This can lead to significant overages for victims, given the scale and size of modern DDoS attacks.

Tip

Check closely and negotiate DDoS protection clauses with your ISP or CSP. Additionally, determine if your ISP uses robust on-prem DDoS protection hardware with a cloud back-up so that small but fast DDoS attacks are mitigated on-prem while large volumetric attacks can be properly mitigated by a cloud DDoS protection service.

MYTH 3

All time-to-mitigate SLAs are created equal

Sometimes numbers can be misleading. Time to mitigate (TTM) is a number often marketed by security vendors. TTM ideally means how quickly malicious DDoS traffic is stopped or blocked, without impacting legitimate traffic and users. It turns out that there's a lot of room for interpretation there. For example, one vendor might not consider a surge in traffic a DDoS attack until it has lasted for at least five consecutive minutes. So the SLA timer may not start until you're already under attack. With the average attack duration being less than five minutes, you can see how this is problematic: It means an advertised 10-second time to mitigate could really be more than five minutes.

Other vendors define time to mitigate as how quickly a mitigation rule can be deployed. It does not reflect stopping the attack or the quality or consistency with which this control is activated. At the end of the day, what you care about is the time to get internet-facing assets secured and back up and running, **with the least amount of impact to legitimate users or services**. Be sure to carefully read the fine print for your vendor's SLA.

Tip

Dig into the details of time to mitigate listed in an SLA. It should represent the equation: The Real Time That Matters = Time to Detect Attack + Time to Apply Mitigation Controls + Time to Block/Stop Attack + Quality/ Consistency of Mitigation. Select a vendor that offers a **true zero-second SLA** for mitigating DDoS attacks without impacting legitimate users.

MYTH 4

Null routing/blackholing and rate limiting are acceptable defenses

Null routing (or blackholing) is a common and rather primitive defensive response from some DDoS mitigation providers. If an asset is under attack and that attack capacity is putting other customers or services at risk, the provider may try to prevent collateral damage by tossing that resource's traffic into a virtual black hole. Does that really help you? From an attacker's perspective, blackholing means mission accomplished — the targeted asset is effectively offline. Depending on the provider's infrastructure, other customers may also end up going offline or experience degraded performance.

Another primitive DDoS defense response offered by many security providers includes putting rate limits on customer traffic as a countermeasure within shared environments. But dropping 20% to 40% of legitimate traffic to give the perception that the asset or service is still up and running is not a successful outcome for the customer under attack. Rate limiting is effective as a secondary or tertiary countermeasure when dealing with DDoS attacks at Layers 3, 4, and 5. When confronting Layer 7 DDoS attacks, rate limiting can be more effective as an initial control but you should always rely on signature mitigation first. You deserve to have 100% of your digital infrastructure effectively protected from DDoS attacks, no matter what layer of the open systems interconnection model it affects, and certainly not just 60% or less.

Tip

Ask your provider how often they blackhole or rate limit traffic during peacetime and when under attack. Determine when (under what conditions) a provider will blackhole traffic and what criteria you'll need to meet to have your services restored.

MYTH 5

It doesn’t matter who shares the cloud platform

Every organization needs security. Controversial businesses that attract frequent attacks, such as gray markets like gambling and adult content websites, need DDoS security defenses, too. Even organizations that promote criminal activity and terror attacks have purchased cybersecurity from legitimate cloud vendors.

It’s easy to think that those sites don’t matter to you. However, if your business shares a cloud platform with an illegal or frequently attacked enterprise, the potential for collateral damage is high. The vendor’s resources may be tied up or overwhelmed, leaving your organization exposed.

Tip

Read a cloud security vendor’s acceptable use policy carefully to confirm that you won’t be sharing security platform resources with high-risk targets. Also, revisit the tips that follow Myth 1 and Myth 2 regarding capacity and capability.

MYTH 6

A web application firewall is sufficient for DDoS protection

Web application firewalls (WAFs), which are often a part of the larger group of web application and API protection (WAAP) solutions, offer effective DDoS protection for application-layer (Layer 7) attacks. While they may offer some basic network layer (Layer 3) or transport layer (Layer 4) protection, it isn’t enough to cover all IPs, ports, and protocols comprehensively.

DDoS attacks come in various flavors and formats, and can target infrastructure layers (Layers 3 and 4), the HTTP(s) application layer (Layer 7), and DNS infrastructure. Moreover, attackers often dynamically switch attacks that could, for example, start with DNS and subsequently expand to other layers or protocols. True DDoS protection comes from a defense-in-depth strategy that adopts a platform of robust solutions that have specific strengths and capabilities to offer protection to Layer 3, Layer 4, Layer 7, and DNS. Any one solution by itself is not always sufficient to cover all bases, and can leave your organization vulnerable to attacks and higher levels of risk for overmitigating legitimate traffic or services.

Tip

Ensure that your DDoS protection solution is not skewed toward one particular type of DDoS attack or implementation design. The best defense comes from a single vendor that can provide multiple dedicated DDoS protection capabilities that maintain interoperability and are supported by a unified rapid response security services team to protect your production resources. The situation becomes complex when these assets are deployed across hybrid networks and cloud-hosted environments. Protection services must be agnostic to the network or deployment model.

MYTH 7

An all-in-one security platform = a better security experience

Some providers offer a variety of services stacked on top of a single-cloud platform. This might reduce the technical complexity of deploying and integrating security controls in the short term, but multiple services that share the same back-end infrastructure and networks are vulnerable to platform outages, collateral damage, and resiliency issues if other parts of the environment are disrupted. Oftentimes, one-stop-shop vendors like this sacrifice feature functionality because of the limitations of their single-platform approach.

A transparent mesh of purpose-built CDN, DNS, and DDoS protection platforms or solutions, designed to solve specific technical and security challenges, means higher-quality mitigation and performance at scale to optimize defensive postures.

Tip

Keep in mind that you don't have to share the same infrastructure to achieve a unified security experience. A defense-in-diversity approach uses underlying architectures that can deliver seamless user experience as well as high-performance security mitigation.

MYTH 8

DDoS protection is not necessary for IPv6

According to Google, roughly 45% of the internet traffic originates from IPv6-compliant devices. In terms of DDoS attacks, IPv6 introduces some improvements over IPv4, such as a larger address space and built-in security features like IPsec, but it doesn't inherently protect against these types of attacks.

DDoS attacks can target both IPv4 and IPv6 networks by overwhelming them with a large volume of traffic, exploiting vulnerabilities, or using various attack vectors that are independent of the IP version. Cybercriminals have already been using the significantly expanded IP space of IPv6 to create even larger volumetric DDoS attacks. In some instances, attackers have sent traffic to random addresses in a network, creating a broadcast storm on the physical network layer and tying up and exhausting router or network resources.

The current fragmentation between IPv4 and IPv6 adds further complexities, since clean IPv6 environments cannot typically be assumed.

Tip

DDoS protection for IPv6 requires similar strategies and technologies as for IPv4, including network monitoring, traffic filtering, rate limiting, and employing specialized DDoS mitigation services.

MYTH 9

You don’t need multiple layers of defense

Most organizations don’t actually believe this myth, but sometimes they build their defense strategy as if it were true. When securing your home, locking your front door doesn’t mean you can leave your back door and windows unlocked. True DDoS defense is achieved by building layers of security that work together seamlessly to prevent attackers from achieving their goal in one single blow.

World-class DDoS defense starts with a network cloud firewall that alleviates the load of your firewalls to the edge of your network. Then a hybrid DDoS protection model will include on-premises, hardware appliance–based protection from short but sharp DDoS attacks, and fall back to dedicated cloud-based protection for large, complex, and volumetric DDoS attacks. Your DNS infrastructure also needs to be protected with a similarly layered strategy that includes using a proxy service that can dynamically implement security policies at the edge of your network and layering it further with an authoritative DNS solution either in primary or secondary mode. Finally, you must protect all your applications and APIs with a robust WAAP solution that includes WAF functionality.

Tip

Layer best-of-breed technologies and solutions with different and dedicated strengths to build a comprehensive defense-in-depth strategy that makes it extremely challenging for cybercriminals to succeed in their attack.

MYTH 10

Every security operations center offers the same level of support

Many vendors advertise security operations center (SOC) support. But having a 24/7 SOC isn't what matters most. What's important is the level of service and expertise you can expect to receive when your assets are under attack. Some key considerations when evaluating DDoS mitigation providers should include:

- What type of support and analysis would you receive before, during, and after an attack?
- How is the SOC staffed to ensure continuity of defense?
- If you contact the SOC, is the person you call the actual analyst performing mitigation, or only the escalation point person?
- Does your provider have security professionals who are trained on mitigation, or are they simply "traffic cops" who route traffic to off-the-shelf mitigation gear?
- Do they offer a custom runbook?

Your security provider's SOC should act as an extension of your incident response team to drive real value.

Tip

Evaluate the expected quality of support you would receive from the service provider's SOC. Aside from attack detection and mitigation, determine if they offer integration and testing, incident troubleshooting, post hoc analysis (lessons learned), and design support to help reduce your attack surface.

MYTH 11

DDoS is an old commodity so the cheapest protection will suffice

The maxim “There is no free lunch” is probably most relevant in DDoS protection. Although a lower price may seem attractive, there are often hidden costs.

Some vendors offer a low sticker price but restrict the number or size of attacks that they’ll mitigate. If you are targeted with too many attacks, or too large an attack, they will ask you to upgrade to a higher (and more expensive) tier of service before stopping the attack — all while you’re trying to get your business back online. Mature DDoS security vendors allow customers the flexibility of choosing between “always-on” and “on-demand” DDoS protection, and switch between them seamlessly, to keep operating costs low while offering best-in-class protection. When comparing vendors and prices, make sure you understand the trade-offs and their impact on your DDoS security posture.

Tip

Understand what’s included in the price you’re quoted before you sign.

DDoS security is complex and it requires significant time and resources within today's rapidly evolving landscape. What worked yesterday might not work today or tomorrow. Staying connected to your end users, customers, and employees is the foundation of your business's success. There's no room for error here — and there's no need to bear the high cost of trying to do it alone. As the most comprehensive, flexible, and trusted DDoS protection platform, Akamai can help.

Learn more about Akamai's DDoS security solutions.

Akamai

About Akamai Security

Akamai Security protects the applications that drive your business at every point of interaction, without compromising performance or customer experience. By leveraging the scale of our global platform and its visibility to threats, we partner with you to prevent, detect, and mitigate threats, so you can build brand trust and deliver on your vision. Learn more about Akamai's cloud computing, security, and content delivery solutions at akamai.com and akamai.com/blog, or follow Akamai Technologies on X, formerly known as Twitter, and LinkedIn. Published 10/24.